CHAPTER

6

SERVER UPGRADES

After reading this chapter and completing the exercises, you will be able to:

- Prepare for a server upgrade
- Adequately test and pilot the server upgrade
- Verify availability of system resources
- Inventory hardware
- Upgrade the processor, memory, BIOS, power supply, UPS, and adapters

Upgrading a server is a more serious matter than upgrading a desktop PC, because it affects a network of users instead of a single individual. This chapter covers the many steps you should take to prepare for the upgrade, as well as how to perform the upgrade itself. It concludes with a checklist to help ensure that you are well prepared for an upgrade and that it goes smoothly.

Preparing for a Server Upgrade

A familiar looking truck pulls up and out hops an overnight delivery employee with a box containing your server upgrade hardware. He places the package on the counter and utters a quick "sign here and have a nice day." Before you even head back to your cubicle, the truck is rumbling down the street for the next "absolutely, positively has to be there overnight" delivery. What you've just seen is the fastest part of a server upgrade.

As the administrator, you must execute every other part of the upgrade with the utmost care and sensibility. First, you probably have to obtain approval to perform the upgrade, which involves budgeting, planning, and presenting your plan to others. Even if you have autonomous power to make decisions for server upgrades, you still need to plan carefully, because if something goes wrong, you are the responsible party. You must also time the upgrade so that it has a minimal impact on productivity and user experience. Be sure to educate users as to what to expect regarding any upgrades that directly affect them. Finally, before performing the physical upgrade, verify that all the server components are there and in working order. With proper planning and precautions, you are now ready to perform an upgrade.

When to Upgrade the Server

On a simple level, an administrator might perform an upgrade in just minutes without notifying anyone. For example, perhaps a server just needs more memory costing less than $200. In many organizations, administrators have discretionary spending for smaller amounts, and it should only take a few minutes to shut down, install the memory, and restart the server. However, upgrading a server can require a significant investment in time, planning, and, of course, money—especially for larger upgrades that have a significant impact on the operations of the IT department, users, or your business audience.

For example, if an organization primarily employs eight NetWare 3.12 servers and would like to upgrade the servers to Windows 2000 Server, it is likely that the new system will need significantly more powerful hardware, as the minimum system requirements for NetWare 3.12 are much lower than for Windows 2000. (Granted, this is a drastic upgrade, but it makes very clear the effect that an upgrade can have on a network.) The organization will likely have to upgrade all major components, including the processors, memory, hard disks, and possibly even the motherboards. A mass migration and upgrade like this would probably be more cost-effective if the servers were replaced with new ones.

Moreover, this upgrade affects more than just how the IT department administers the servers—network resources for users might be located differently. Logon scripts mapping to NetWare printers won't work anymore, and users that have manually created

shortcuts to network resources will no longer be able to access those resources. The impact of an upgrade this drastic would require significant planning, proper timing, a smoothly executed upgrade process, adequate personnel, and plenty of communications with the users notifying them of the upgrade and its impact on their day-to-day functions.

TIP

A migration from NetWare to Windows 2000 servers can be less disruptive if you retain the NetWare server names. That way, shared resources will have the same network path if you also keep the same directory structure. Also, consider gradually upgrading one server at a time using a pilot program (as described later in this chapter), and check the impact as you go. This might help prevent a system-wide network failure or disgruntled users if something goes wrong.

Many times, upgrading the server is a necessary step in response to poor server performance. To fully justify expenditures in time and money, you should create a performance **baseline** so that you can define an acceptable level of performance. (Chapter 11 discusses performance monitoring and optimization, including the establishment of a baseline.) If you see the server begins to perform poorly on a regular basis when compared to the baseline, consider upgrading components that might be a bottleneck. For example, a heavily utilized database server will have higher demands placed on the processor than a simple file server. Upgrading or adding another processor might help the system to perform better.

Caution

Before upgrading a server component, pause to verify identification of the actual bottleneck. For example, a server that shows heavy hard disk utilization might not need a larger or faster hard disk. Instead, first check memory utilization. Most NOSs utilize a swap file mechanism that substitutes hard disk space for memory in a low memory situation. Therefore, heavy hard disk utilization could actually be a memory issue that would be minimized with a memory upgrade.

You might take that same database server and consider upgrading it not only to an acceptable level of performance, but beyond current needs in a proactive approach to extend the server investment. This might prevent the repeated expense of upgrading again in a few months as demand on the server grows. For example, if the database server has one 700 MHz Pentium Xeon processor now, you might consider installing two 1 GHz processors, if the motherboard supports it. A proactive upgrade such as this is sometimes also in anticipation of financial timing. For example, in some organizations, if the department does not use its entire budget in a given time, then the next budget allocation might shrink because of a perception that the department does not need as large a budget since it didn't spend the entire budget last time. Although this kind of budgeting model has obvious weaknesses in logic and wisdom, it is nevertheless a reality in many organizations.

Other reasons for a proactive approach might include the anticipation of an upcoming merger with a new parent company. Perhaps your current intranet is for general information only and does not have a high hit rate. However, the new parent company requires you to post a great deal of company information on the intranet, such as details on employee benefits, company announcements, web-based collaboration software, and so forth. In this case, you might consider several upgrades, including adding another NIC to the server, to increase network throughput in anticipation of a higher hit rate.

Timing the Upgrade

Generally, the IT professional tries to keep as low a profile as possible within his or her organization. Why? Because when the administrator has a high profile, it's usually because something is wrong with a server or the network, and users tend to immediately think of administrators as having caused the problem! When users don't notice your presence in the organization, it usually means you are doing your job well. An ill-timed server upgrade (or scheduled maintenance) can give the administrator a very high profile; therefore, you should make sure that days or weeks before the upgrade, you monitor server utilization to determine periods of peak usage and periods of lowest usage. Of course, you want to perform any upgrades that make the server unavailable only during periods of low usage. If your organization is most active during business hours, you can stay late or arrive early to perform your upgrades while nobody else wants to access server resources. However, many organizations, especially larger ones, operate continuously. In addition, if the organization is global, users on one side of the globe might access resources while users on the other side of the globe are sleeping. This type of continuous access means that some types of upgrades will definitely take place while some users are attempting to gain access to resources. If the server must be taken offline, you can either temporarily transfer the role of server to another server or you can notify users in advance so they don't expect access during the time of the upgrade. Some upgrades (such as adding a disk to an array) might not require any downtime, and users will never notice the difference. On the other hand, upgrades such as adding processors or memory require the server to be turned off, and you should schedule an appropriate time and notify users in advance.

Notifying Users

Notifying users of a server upgrade helps to reduce the administrator's visibility and avoid unnecessary calls to the IT department. Even if you try to upgrade the server during the lowest usage periods, someone will still wonder why the server is unavailable, so the notification should help. Also, a public advance notice such as an email broadcast shows that you made a reasonable effort to notify users.

Broadcast company email is a common notification method, as are notices on the company Internet or intranet. Start notification as far in advance as is practical, which will vary from one organization to the next. (Some well-organized organizations have a written administrative policy for planned downtime notification.) I recommend an initial

notification far in advance, and then as the upgrade approaches, notify a few more times with increasing frequency up until the actual upgrade event.

Note: Notification is not important *only* for when users lose access to the server. If the server is a load-balancing server, users will notice slower response and you should notify them in advance.

Notify the users when the server will be unavailable and for how long. If possible, also state how the planned downtime benefits the user. This helps to psychologically cushion the inconvenience for users when they understand it is ultimately for their benefit. For example, most users would be grateful if you were to add more storage to the email server so they could store more messages.

TIP: When a server goes offline unexpectedly and not as a result of upgrades or maintenance, you should attempt to send out a message to all users to stem the certain flood of calls about the inaccessible server.

TIP: Try to maximize the upgrade process so that other tasks can be completed in the same approximate time without increasing the impact on user access. For example, any time you want to work inside the server case, you might as well get a vacuum or a can of compressed air and eliminate the dust. Since that's probably also a regularly scheduled maintenance item, doing it now saves the separate task of doing it later.

Confirm That You Have Necessary Upgrade Components

Avoid unproductive downtime during an upgrade by confirming that all the necessary components for a successful upgrade are accounted for. For example, some hardware might not respond appropriately to the operating system without the proper BIOS upgrade, so you need both the BIOS upgrade and the hardware itself. When you perform an upgrade, be careful about the drivers that come with the hardware. (A **driver** is a software interface that allows the hardware to function with the operating system.)

You do not know how long the device has been on the shelf, and the drivers might be outdated. Save yourself the task of time-consuming troubleshooting after the upgrade that might occur due to an outdated or incompatible driver. After checking to see if the device is compatible with the other hardware and network operating system (NOS), download the most recent driver from the vendor's web site. While you're there, check the FAQ section to address and prepare for any issues you might encounter during the upgrade.

TIP: Download and expand the drivers into a permanent network directory from which you install the drivers. You might need to reinstall the drivers from time to time, and many NOSs default to installing drivers from the original installation path. Also, the drivers are immediately available to other servers on which you perform the same upgrade. Make sure you dedicate each directory to only a specific vendor and a specific device, because driver files might have the same name (especially from the same vendor) and you might accidentally overwrite files.

Similarly, check with the NOS vendor to see if there are known problems with the particular device and/or driver you want to install. A Readme.txt file often accompanies drivers and updates. Although this file is usually a statement of obvious information ("this driver upgrades your network card"), it might also contain important information ("this update only applies to Windows 2000"). Many administrators ignore this little file, but at the very least, you should scan it for any red flags or installation tips. The NOS web site often informs you of incompatible devices and provides solutions. Often, the solution is to avoid specific conflicting hardware devices or to install a NOS upgrade, patch, or hot fix. If the NOS vendor's support team does not list any known problems, check with other sources, such as newsgroups focused on the specific hardware and/or NOS vendors.

Caution: We administrators tend to be somewhat stubborn about reading directions, preferring to figure things out for ourselves. However, it often pays to read instructions in order to avoid problems later. I recently installed a security camera that attached to the USB port of a server, hoping that the Windows 2000 Server Plug and Play detection process would find the drivers on the camera's CD-ROM, as with most other installations for which Windows does not have a driver. Nothing worked. Had I read the instructions, I would have known that you must install the software and drivers first, and then install the camera. This would have saved me quite some time in trying to troubleshoot something that had a plain solution to begin with.

The level of precaution recommended in these pre-upgrade tasks might seem overly cautious. If you were performing a simple upgrade on a home PC, it might be. However, because of the impact the server has on an organization, every precaution is necessary.

Planning for Failure

Anybody who has administered computers for more than a few weeks knows that even the most meticulous and professional upgrade attempt sometimes encounters unanticipated problems, conflicts, or incompatibilities. In preparing for this possible contingency, administrators should prepare adequate failsafe measures to quickly recover from the problem or continue service to users through redundancy while troubleshooting the failed upgrade.

Always back up the server before performing any hardware or software upgrades. I recommend that you do not depend upon the normally scheduled backup rotation, because if a problem occurs, the backup can be slightly outdated. For example, most backups take place in the middle of the night, but if you perform the upgrade after everyone leaves work but before the backup, then one day's working data is in jeopardy if a problem occurs. Instead, take the server off of the network so that new data cannot be written to it. Next, perform a full backup of at least the data, and possibly also the operating system. Now you have a snapshot of the system before the upgrade, and if necessary, a restore should replace the data intact.

TIP

Consider using imaging software such as PowerQuest DriveImage or Norton Ghost. This exactly duplicates the hard disk to a single image file for restore should a problem occur. You can store the image on a network share or one or more CD-ROMs, and protect it with a password in case the CD-ROM falls into unauthorized hands. Alternatively, many network cards are now bootable, allowing you to access the network and restore the image to the server from the network share. If you have to restore the image, it is much faster than reinstalling the operating system, reinstalling all of the server applications, and then restoring the data from a conventional backup.

In addition to creating a backup, other failsafe methods can also ensure server availability. For example, recall from Chapter 1 that clustering is utilization of two or more servers hosting the same application. If one of the servers is unavailable (as might occur during an upgrade), the remaining servers in the cluster continue to provide service. With a mission-critical server or application, you probably already have clustering enabled.

Absent a cluster, administrators might have a **hot spare**, which is a specific component (usually a hard drive) or a complete server that can immediately be available on the network and transparently perform the exact same functions as the original.

Verifying System Resources

You never seem to have enough PCI slots for the devices you want to install in a server. While you might have a free ISA slot, you probably don't care because fewer devices are ISA compatible. However, PCI slots represent valuable slot real estate that quickly fills up. In a typical server, you probably have five PCI slots. Account for two network cards for better throughput and availability, a SCSI card for tape devices, and another for hard disks, and you're almost out of expansion slots already. If so equipped, you can add a mezzanine or riser board to the motherboard to expand the number of available slots. Whatever the case, in larger environments the administrator probably does not know offhand exactly how many expansion slots remain in each server. Before making plans to add a device, verify that sufficient slots are available.

Even if a slot is available, you might encounter issues with available IRQs, DMAs, or I/O ports. These are each limited resources that most devices require to communicate with the operating system and other devices, and they are defined as follows:

- An **interrupt request (IRQ)** is a request that the device uses to "interrupt" the processor to ask for processor resources. There are 16 IRQs, numbered 0–15. Several IRQs are preassigned. For example, the COM1 serial port usually has IRQ 4. PCI IRQ steering (described in Chapter 3) can allow multiple devices to use the same IRQ if no more unique IRQs are available. However, ISA devices cannot take advantage of this benefit.
- A **Direct Memory Address (DMA)** is a resource that ISA devices use to directly access memory without first having to access the processor, both increasing device performance and reducing processor load. There are eight DMA channels, numbered 0–7.
- An **I/O port** is a location in memory that the processor uses to communicate with a device.
- A **memory address** is a dedicated region in system memory that some devices reserve and that is unavailable for use by any other device, application, or the operating system. This can help device stability by ensuring that nothing else trespasses the memory, which causes system errors.

If you are out of IRQs and IRQ steering is not available, or if you are out of available DMAs, then you must remove (or disable) an existing device that requires those same resources, or you cannot upgrade the server. I/O ports are usually plentiful, and if two devices request the same I/O port, you can usually reassign one of them to an alternate port.

Depending upon the chassis and power supply, you might not have sufficient expansion space to add more hard disks, tape drives, or removable storage such as Zip drives, CD-ROMs, or DVD devices, all of which require a drive bay, either internal or external. Drives installed in an internal drive bay are neither accessible nor visible when the case is attached. Most commonly, you install hard disks internally. It would be impractical to install removable storage or CD-ROM/DVD drives internally, so you want to use an external drive opening, which means that you can access and see the drive. Besides drive bay availability, you need a power supply that can handle the additional power requirements of the devices and has sufficient power connectors for each drive. You can use a Y-cable split that converts a single power connector to two, but if you're using several of these, you might be overloading the power supply. Generally, the more powerful the power supply, the more power connectors it includes.

> **TIP** If you want to install a 3-1/2 inch hard disk but only have a 5-1/4 inch drive opening, purchase an inexpensive adapter kit that adds brackets to the outside of the drive, expanding it to fit into the larger 5-1/4 inch opening. Sometimes the adapter kit is included with the new drive.

A standard desktop or entry-level server tower case has space for perhaps four internal drives and two or three external drives. A server has significantly more storage space. For example, the Compaq ProLiant 8000 has internal drive cages for 21 hot-plug hard disks. Regardless, verify available drive bays as necessary.

Caution

If you plan to install a large quantity of internal hard disks, consider adding one or more additional cooling fans to compensate for the additional heat.

Making an Inventory

One of the most frustrating things about installing hardware is finding it! Even relatively small organizations quickly accumulate quite a few loose components, chips, hard disks, and so forth. When a server fails, you must know exactly where to find replacement parts, so you should carefully inventory (and lock up) all parts that have any value. Organize smaller loose parts in appropriately sized storage bins, trays, and cabinets. Having an inventory also helps you to control and be aware of possible theft.

In addition, you should know what equipment is in each server for proper asset tracking, budgetary projections, and warranty service. Especially in large environments, manual inventory of installed hardware is an arduous and seemingly endless task, further complicated when there are multiple sites. I recommend procuring software that can automatically scan your entire network to inventory not only installed hardware in your servers and clients, but also installed software. Some programs can also identify network devices such as hubs, routers, switches, and so forth. One of the most popular programs is Microsoft Systems Management Server, for which you can find detailed inventory instructions at *www.microsoft.com/technet/SMS/c0318341.asp* or perform a search for the title "Administering Inventory Collection." Also check into the following other vendors:

- Hewlett Packard's OpenView at *www.openview.com*
- IBM Tivoli at *www.tivoli.com*
- Computer Associates Unicenter TNG at *www.cai.com*

Note that all these products require hardware that is capable of responding to queries from the software. You can still manually inventory hardware that does not automatically respond to the software, but this is becoming less of an issue as more hardware is designed to be compatible with inventory software.

When receiving new equipment or equipment transferred from another office, always request that an inventory list be included with the shipment. This helps to ensure that equipment arrives as promised and that the server from the home office that has 512 MB RAM doesn't suddenly appear in your office with only 128 MB. Another reason for the inventory list is that it helps you in assembling the equipment. Many servers and their associated equipment involve dozens of parts, including zip ties, cable management systems,

fans, screws, keys, books, warranty cards, power cords, and so forth. It is extremely frustrating to unpack and assemble an entire server and rack, only to find that you are missing a vital component that you would have known about had you compared the physical parts to an inventory list.

TEST AND PILOT

The potential impact of some upgrades (both if they succeed and if they fail) might require an isolated **pilot program** in which you thoroughly test the upgrade for reliability and performance prior to deployment throughout the organization. A pilot program isolates a server upgrade in a portion of the network that makes performance easier to determine and lessens negative impact should some part of a major upgrade fail. Some organizations require a pilot program for nearly any hardware, operating system, or software change. However, pilot programs are not usually intended for common upgrade items such as installing a hard disk, network card, or memory. Most implementations might involve something like an upgraded NOS, an entirely new NOS (migrating from UNIX to NetWare 5.x, for example), or a change in hardware architecture (such as a change from UNIX-based Alpha architecture to Windows-based Intel architecture). If the upgrade works well in its initial pilot, you can extend the pilot programs to other segments of the network to see if the results are also successful under different circumstances. Finally, when the upgrade is fully tested and has satisfactorily passed the pilot phase, you can deploy the upgrade throughout the remainder of the organization as necessary. As an example for the current context, which focuses mostly on hardware, consider a pilot program for upgrades of the BIOS, motherboard, processor (especially if changing platforms from, say, AMD to Intel), and anything else you think might significantly impact the network if it fails or requires isolated analysis.

TIP The CompTIA Server+ Exam Blueprint recommends pilot programs when upgrading processors, hard disks, memory, BIOS, adapters, peripherals (both internal and external), service tools, and the UPS. This doesn't leave much out. Though in the real world you might not pilot every item (who implements a pilot program for a new mouse?), approach the exam as if you would.

Note Sometimes, a test lab environment precedes the pilot phase. The test lab can be a special network segment completely detached from the remainder of the network in which administrators can perform drastic tests on the server without concern for affecting users or other production servers.

Once you successfully implement the upgrade changes, make sure that the improvement is more than just a perception. Start recording the performance of the server and/or network as it applies, and compare it to the previous baseline and performance prior to the upgrade. You should be able to find an improvement in the targeted upgrade area. For

example, if you added another processor, you should see overall processor utilization drop to a lower percentage (where lower percentages equal better performance). After recording performance statistics, you should be able to change the baseline's level of acceptable performance. In most organizations, the baseline is a moving target that you periodically reset as server demands increase. If responsible parties (you and management) determine that network or server responsiveness no longer meets an acceptable range of performance, then a decision must be made: Either reset the baseline at the new level of performance or modify network or server equipment to return to the original (or better) level of baseline performance.

> Note
> When confirming a successful upgrade, you can check for the obvious items such as proper functionality of the hardware. Don't forget, however, that most NOSs have logs that might also record errors in a problematic upgrade.

Be sure that after performing the upgrade, you record it to an easily accessible source for troubleshooting and asset-tracking purposes. Some organizations might have a log book next to the rack where changes are hand-written; others might have a computer-based log, such as a database or a spreadsheet, saved to an administrative network share. I prefer the latter method, because the log is accessible from any administrator's desktop. For example, in troubleshooting a server in Denver, an administrator in Phoenix can open the server log on a network share and see if any recent hardware upgrades might have caused a problem.

PERFORMING THE UPGRADE

When you've completed all the steps outlined here to prepare for an upgrade, you should be ready for the hands-on work of upgrading a server or server components. Before you start, though, be sure that you've taken precautions to prevent electrostatic discharge.

Avoiding Electrostatic Discharge

Before touching anything inside the server, it is critical to exercise precautions against **electrostatic discharge (ESD)**. You have probably experienced ESD at no significant harm to yourself many times, particularly if you live in an area where winters are cold and the wind blows (Chicago is a great example!). Once you are in from the cold and touch a door handle (or pet the cat), an ESD occurs. ESD occurs when two objects with differing electrical potential come into contact with one another because the electrical charges seek to equalize. While you are outside in the cold wind, the energy from the wind can build up electrical potential in you. When you touch an object in the house with less electrical potential, static electricity discharges from you to that object.

Although people are hearty enough to sustain a static shock, servers and their components are not. Before you touch a server for any reason, including component installation or inspection, you must be certain that you present no ESD threat.

Caution

Although it doesn't seem like much, if you can feel the static discharge, then you probably discharged around 3500 volts. If you also hear the discharge, then you probably discharged around 5000 volts, and if you see the discharge in a lighted room, then you probably discharged around 8000 volts. It is easy to build up this voltage—walking across the floor can generate 15,000 volts, and removing bubble pack from a carton can generate as much as 26,000 volts. It only takes 100–1000 volts to negatively impact server components, and at lower levels you might not feel, hear, or see the discharge.

You might unknowingly damage a component with ESD in two ways. First, an upset failure affects only the reliability and/or performance of a component. This is perhaps the worst of the two types of damage because it is difficult to detect and consistently reproduce. (For example, it takes 200–3000 volts to damage a server's CMOS.) The second way ESD affects a component is a catastrophic failure, which immediately damages the component so that it ceases to function properly.

Note

Semiconductor Reliability News attributes approximately 60% of electronic component failures to ESD.

In seeking to prevent static discharge, take the following precautions:

- *Touch the chassis*—Touching the chassis grounds you and equalizes the voltage levels between you and the server. This is not a sufficient precaution, however, because as time passes, voltage can build up in you again, particularly if you are wearing leather soles on carpet. Touching the chassis is only a temporary, initial precaution. If you cannot utilize a better method, such as using an ESD protection kit (see below), then continue to touch the chassis periodically.
- *Unplug the power*—Many people assume that because the power cord is plugged into a grounded outlet, the case is protected against ESD. While it is true that the plug does lead to earth ground, that is not what is important in preventing ESD. What you are seeking is equalization in electrical potential. Plugged or unplugged, touching the chassis temporarily equalizes electrical potential. Moreover, it is safer to unplug the server. By leaving the server plugged in, you risk accidentally bumping the power switch. Installing or removing a device with the power on is catastrophic to most components (unless they are hot-swappable) and introduces the risk of system-wide electrical damage, not to mention the unpleasant surprise of receiving an electric shock. Also, many power supplies continue to supply low-voltage power to

the motherboard, even when switched off. By implementing proper grounding measures and unplugging the server, you can avoid a potential mishap.

- *Use a grounding kit*—These come in several forms. At the lowest end, a portable grounding kit uses a wrist strap with an alligator clip that attaches to the server chassis (see Figure 6-1). (Be sure to attach it to an unpainted surface for best contact.) This has the same effect as touching the chassis to equalize electrical potential, except that it is not temporary because the connection is constant.

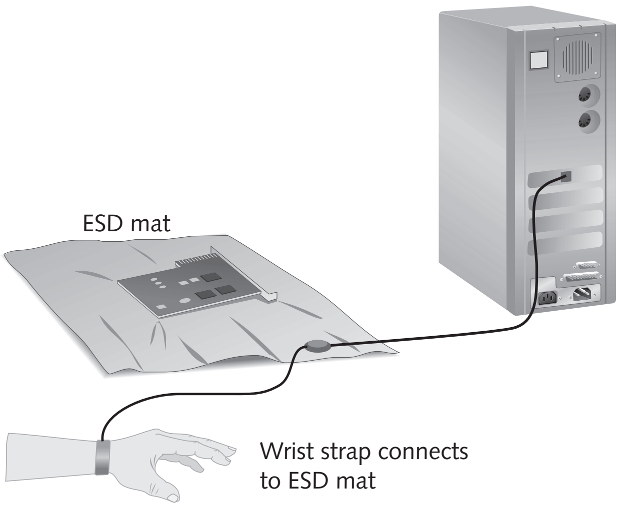

Figure 6-1 A portable grounding kit and mat

A hardware repair bench normally includes a grounded floor mat (also called a map) that connects to earth ground at a nearby electrical outlet (this does not affect the operation of other devices plugged into that outlet). Another mat on the benchtop is also grounded—either to the floor mat or independently to another electrical outlet. The user wrist strap connects to either mat or the server chassis (see Figure 6-2). Because both mats, the wrist strap, and the chassis are all grounded, they possess the same electrical potential, eliminating the risk of ESD. The wrist strap usually includes a resistor designed to negate a high-voltage electrical charge, in case the technician accidentally touches a high-voltage item such as internal components in the power supply or the monitor. Both items retain high amounts of voltage even when unplugged. Full-time repair facilities normally ground the entire workbench.

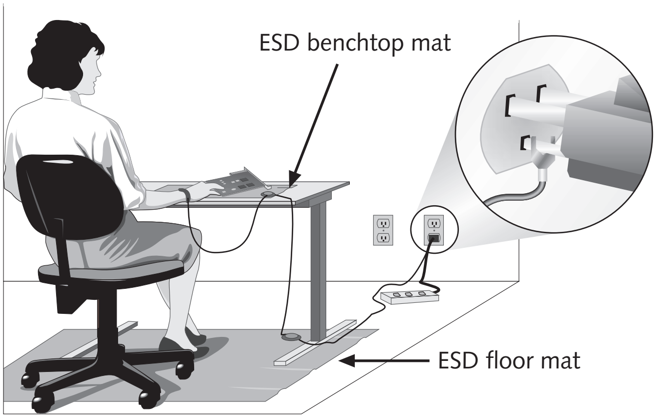

Figure 6-2 A workstation grounding kit

For more information on specific grounding kits, refer to the following web sites (also see *www.3m.com/ehpd/esd_training* for more about ESD in general):

- Specialized Products Company at *www.specialized.net*
- 3M Corporation at *www.3m.com/ehpd/workstation*
- Jensen Tools at *www.jensentools.com*

TIP

If you find that the cord attached to the wrist strap gets in your way as you work, you can instead use a heel strap that attaches to your shoe.

Other tips for avoiding ESD include:

- *Take it to the mat*—As you install and remove components, be sure to place them on the mat. You might be tempted to place the computer itself on the mat, but that is unnecessary, provided you grounded it properly. Leave it off the mat so that you have some workspace.

Caution

Some people recommend placing components on a sheet of aluminum foil. I cannot recommend this because some components contain tiny built-in batteries. If the batteries short out, they can become extremely hot in an instant, and might even explode (think firecracker). Batteries are not always easily identifiable on the board.

- *Handle with care*—If you find yourself without ESD protection, be sure you handle loose adapter cards by the metal bracket that attaches to the chassis. The internal ground circuitry of the card is connected to the bracket, so touching the bracket prevents ESD from damaging the card components. If the device does not have a bracket (a motherboard, for example), handle it by

the edges and try to avoid touching any of the items on the surface of the card. Do not touch the metal edge of the adapter that goes inside the expansion slot, because even minor soiling from oils in your hand can contribute to corrosion. Do not stack components.

- *Bag it*—Place loose components into ESD-resistant bags. Remove any other type of packing material, such as Styrofoam, bubble wrap, or cellophane, because they tend to build and hold a static charge.
- *Take off the jewelry*—Metal jewelry can conduct electricity. It is a good idea to remove any jewelry from your hands and wrists (watch, rings, bracelets) before working on the server. Also, jewelry can catch onto components and wiring and hinder your dexterity.

Now that you've made all the preparations for a server upgrade and have ensured that you are protected against ESD, you can proceed to upgrade server components.

Upgrading the Processor

Before upgrading the processor, perform a few tasks to prevent serious problems or damage. Primarily, you want to verify that the BIOS and motherboard support the processor, and if using SMP, that the new processor is compatible with the existing processor.

Generally, it's a good idea to keep your BIOS version as up-to-date as possible—this chapter addresses this issue later. Because the BIOS directly affects the communication between the processor and the rest of the system, the importance of BIOS compatibility is obvious. Besides keeping up-to-date, also access the BIOS settings and verify that it supports SMP. A PC server usually allows 4-way SMP unless a mezzanine board or other motherboard modification allows you to expand to 8-way or greater SMP.

Next, check the motherboard to see if the proposed processor is compatible. The form factor of a given processor might physically fit in several different motherboards, but that does not mean it is compatible. Recall that a processor operates at multiples of the bus speed. This is one of the factors that limits the available upgrade path in your server. For example, you cannot replace a 700 MHz Pentium III with a 900 MHz Pentium III. The 700 MHz Pentium III is designed to operate on a 100 MHz bus. However, the 900 MHz processor is designed to operate on a 133 MHz bus, which will not allow the processor to function properly on a 100 MHz bus. Also, a given chipset might not be compatible with the proposed processor. To determine the compatible processor upgrade path for the server, you could research the motherboard manufacturer's web site. However, it is better to verify the upgrade path with the actual server vendor because they might have integrated something else into the system that affects upgrade compatibility.

Adding another processor to an existing processor involves more than simply making sure that both processors are the same speed. The processors should be identical in every way, including cache size, form factor, and stepping. For example, there are at least eight different 700 MHz Pentium III Xeon processors. While they all operate at 100 MHz,

they vary in L2 cache size—either 1024 KB or 2048 KB—and the new processor cache must match the existing processor cache size. When adding another processor, verify that the new processor's **stepping** (the processor version) matches that of the existing processor. As Intel manufactures processors, minor problems, incompatibilities, or inaccuracies might be discovered from time to time. While the chances that these flaws will negatively affect server operation or compatibility are minimal, Intel usually corrects them when practical, so two processors of the same speed and cache still might not be exactly the same. You can also look on the processor to find its specification number (or **S-spec**)—an alphanumeric code that uniquely identifies each processor version and is more specific than the processor stepping. Figure 6-3 shows an S-spec (the last item on line 3, SL4MF). Notice also that the first line indicates 1000 MHz, 256 KB L2 cache, 133 MHz bus speed, and 1.7 V power.

Figure 6-3 The S-spec on this processor is SL4MF

Processor Slots and Sockets

The processor and its respective slot or socket appear in two primary formats. Sockets accommodate a processor format known as the **Pin Grid Array (PGA)** processor, which can be a flat, thin ceramic device with hundreds of gold pins on the bottom. These types of processors fit inside a motherboard socket receptacle that accepts each of the pins on the processor. The most recent implementation is the **Staggered Pin Grid Array (SPGA)**, which staggers the pin arrangement to squeeze more pins into the same space (see Figure 6-4). The PGA has two formats: standard PGA and the flipped chip PGA (FCPGA), referring to the fact that the processor die is "flipped" upside down on the die. The PGA format makes no difference in the actual installation except that FCPGA includes a fan in addition to the heat sink (described later in this section).

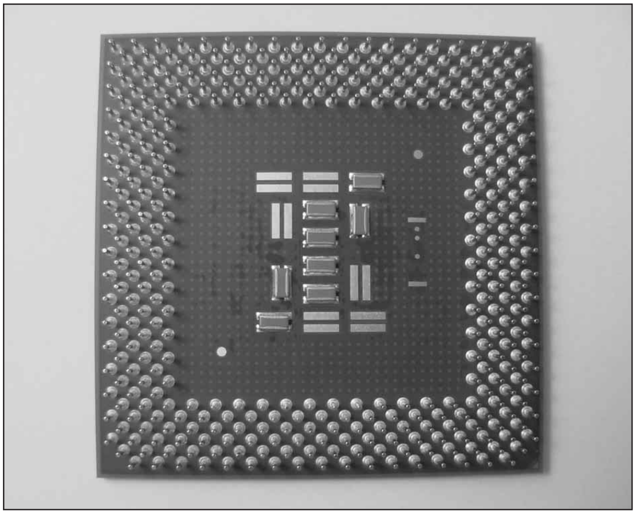

Figure 6-4 An SPGA processor

The slot format processor can be a larger device that stands upright inside a motherboard slot, similar to adapter or memory slots (see Figure 6-5). This format is referred to as the **Single Edge Contact Cartridge (SECC)**. The **Single Edge Contact Cartridge2 (SECC2)** format is a similar form factor, except it exposes the contacts at the bottom. The slot format processor includes a specially constructed plastic and metal housing that often includes an on-board cooling fan. Processor manufacturers have flip-flopped over the years regarding which format they use. Most processors use the socket format; however, Intel seems to favor the SECC slot format for the Xeon at this point. Whichever format is used, we will consider the housing and CPU collectively to be the processor.

Figure 6-5 A Pentium III Xeon processor in the SECC format

Table 6-1 lists the types of sockets and slots in current use.

Table 6-1 Sockets and Slots Used by Various Processors

Socket or Slot	Processor
Socket 370 (or PGA370)	Socket versions of the Intel Pentium III and Celeron
Socket 7 (or Super 7 when faster than 66 MHz)	Intel Pentium, Pentium MMX, AMD K5, K6, K6-2, K6-3
Socket 8	Pentium Pro
Socket A	AMD Duron and PGA format Athlon
Slot 1 (or SC-242)	Slot versions of the Intel Pentium III, Celeron, and Pentium II
Slot 2 (or SC-330)	Intel Pentium II and III Xeon
Slot A	AMD Athlon

Inserting the processor, either AMD or Intel, into a socket is an easy matter. Look carefully at the pins on the processor and match them to the socket on the motherboard. You will usually see a bevel that prevents you from accidentally inserting the processor in the wrong orientation. Simply match the bevel on the processor to the bevel on the socket. Before you insert the processor, lift up a lever next to the socket. The lever is the lock that holds the processor in place, and the feature is known as **zero insertion force (ZIF)** because when you insert the processor, gravity alone should be enough to seat the processor into the socket (see Figure 6-6). Sometimes you might have to help gravity a little bit, but very little pressure is required. If the processor does not seem to drop easily into the socket, do not force it. If you do, you may be buying another new processor.

Slots are keyed so that, again, you can only insert the processor in one orientation. However, be aware that it requires significantly more force to insert the processor into a slot, and the retention mechanism has guiding slots that facilitate this (see Figure 6-7). Use the retention mechanisms on either end of the processor to release it from the slot if you need to remove it later, but sometimes that can be a challenge as well. Most current processors require a slot.

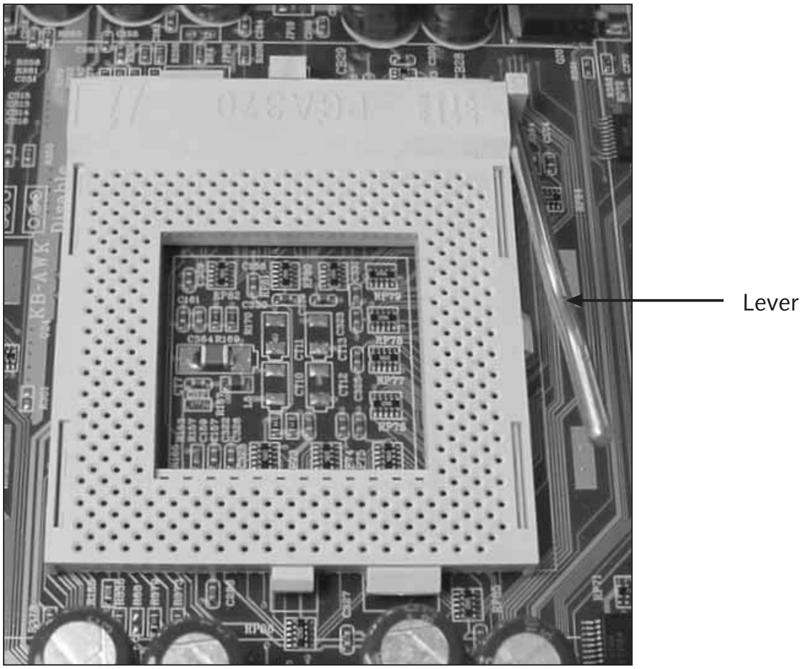

Figure 6-6 A ZIF socket uses a lever to lock the processor in place

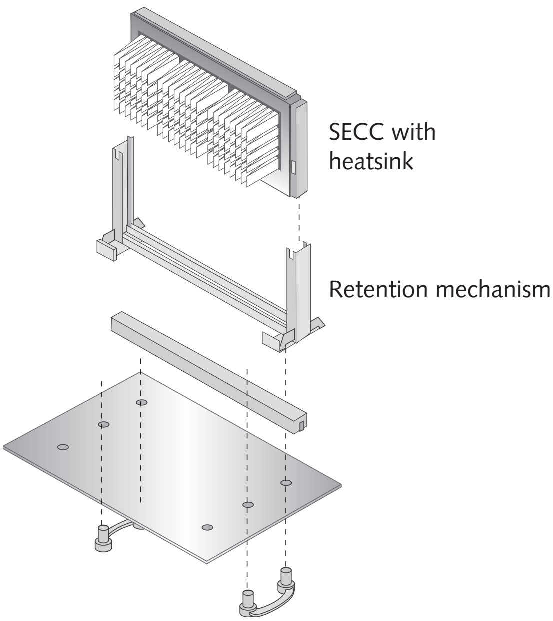

Figure 6-7 The slot and its retention mechanism

Removing the slotted processor properly is a matter of experience and getting a feel for removing the processor, which is usually seated very firmly in the slot. You have to balance the objective of removing the processor with a prudent degree of restraint, and this might require some experience. A safer method is to procure a processor extraction tool. Flotron (*www.flotron.com*) makes such a tool (see Figure 6-8).

Figure 6-8 A processor extraction tool

Processor Cooling

Processors can get very hot, and you must exercise care in cooling them properly using a **heat sink**—an attachment to the processor that either dissipates heat passively, through aluminum cooling fins, or actively, using a small cooling fan, usually in addition to cooling fins. When you purchase a "boxed" processor, it includes a cooling solution from the manufacturer. If it is "bare," then it has no cooling solution and you must determine a way to cool it yourself. Processors with only passive cooling depend upon airflow from the power supply and/or other cooling fans in the system. I usually recommend using active cooling even when only passive is required and even if the boxed processor only includes passive cooling.

Other server components, particularly ribbon cables for hard drives, can block airflow. Be sure to route cables so as not to impede airflow.

Socket processors usually have a fan mounted on top of the cooling fins. The short fan power cable is usually sufficient to reach the motherboard power connection for the CPU, and is often marked on the motherboard as "CPU FAN." Attaching this type of

heat sink to the CPU (see Figure 6-9) typically uses a clip that you hook to a notch on one side of the socket. On the other side, you press down on the clip until it hooks onto the notch on the other side of the socket. Figure 6-9 shows the aluminum cooling fins, clip, fan, and thermal tape.

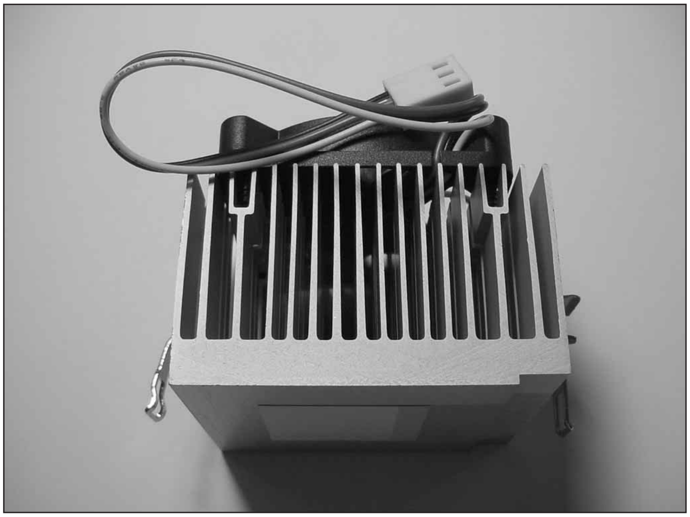

Figure 6-9 A heat sink

Figure 6-10 shows the installed heat sink, fan, and fan power connection to the motherboard. Inserting the heat sink might take considerable force; take care that the heat sink is oriented and aligned properly.

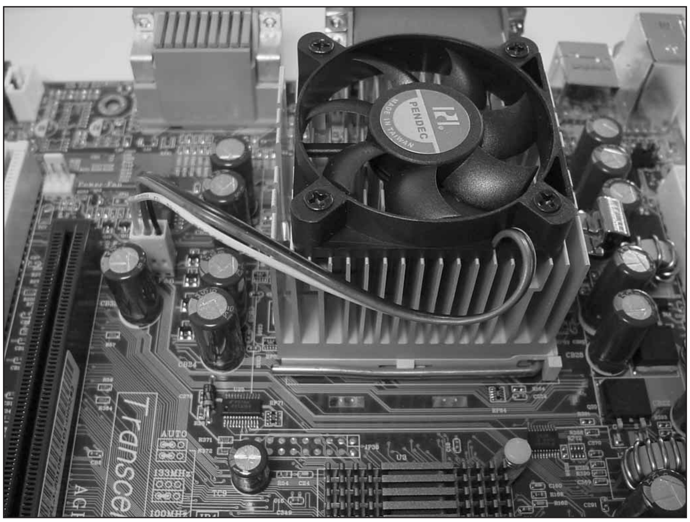

Figure 6-10 The installed heat sink, fan, and fan power connection

If you want to provide maximum cooling, you can buy third-party heat sinks and fans that are usually much larger than those that come from the manufacturer and provide even better cooling. Normally, these are no more than $50.

Caution

Be careful that other cables do not contact the active heat sink cooling fan, because the resistance from the contact slows down the RPMs (and hence the cooling) and shortens the life of the fan.

Processors that go into slots, such as the SECC for the Pentium II/III or the SECC2 for the Pentium III Xeon processor, use a similar cooling method to the socketed processors, except that the form factor is rectangular and larger to accommodate the larger slotted processors.

TIP

Most heat sinks include a small amount of thermal tape located at the contact point between the bottom of the heat sink and the surface of the processor. The purpose of the thermal tape is to act as a conductor through which heat is transferred from the processor to the heat sink. Otherwise, there would be a narrow gap of air between the processor and heat sink, and air by itself is not a good conductor. Many technicians prefer to apply inexpensive thermal grease instead. A small amount of grease fills the gap and draws away heat better than thermal tape.

Notifying the Operating System

In Windows NT and Windows 2000, you must notify the operating system of a change from one processor to multiple processors; otherwise, the operating system does not recognize or use the additional processor(s). The steps to make this change in Windows 2000 are as follows:

1. Right-click My Computer and choose Manage.
2. Select Device Manager in the left tree pane.
3. Expand the Computer node in the right tree pane. The type of computer appears under Computer, and is probably Standard PC or Advanced Configuration and Power Interface (ACPI) PC.
4. Whichever type of computer appears, right-click it and choose Properties.
5. Click the Driver tab of the Properties sheet.
6. Click the Update Driver button.
7. Click Next to skip the introduction of the Upgrade Device Driver Wizard.
8. Select Display a list of the known drivers for this device so that I can choose a specific driver, and then click Next.

9. Select the Show all hardware of this device class option. The screen shown in Figure 6-11 appears.
10. Select the multiprocessor option that matches your computer, and proceed to the end of the wizard.

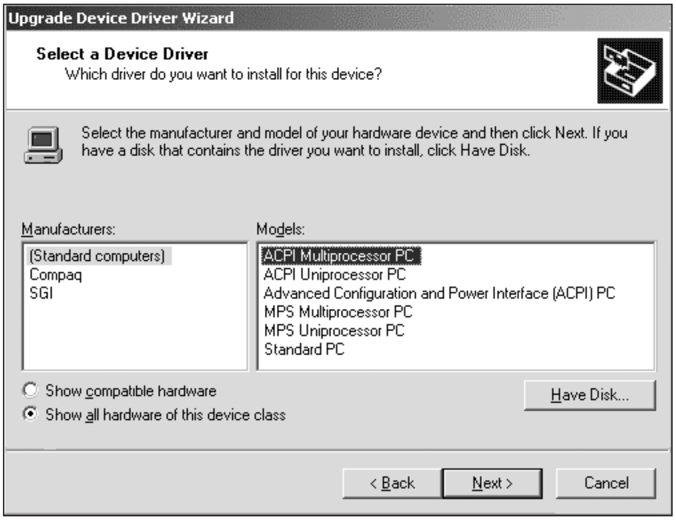

Figure 6-11 Choose ACPI Multiprocessor PC or MPS Multiprocessor PC as it applies

Microsoft offers a series of utilities with the Windows NT 4.0 Resource Kit. Run the UPTOMP utility from the command prompt, or reinstall the operating system to use multiple processors.

Note In current servers, it is uncommon to have to set a jumper to the correct voltage, CPU multiplier, and bus speed, but check documentation first, especially for Socket 7 or Super 7. If the voltage is set improperly, the processor might not function correctly or excessive voltage might damage the processor. Socket 370, Slot 1, Slot 2, and Slot A each adjust the voltage automatically. However, BIOS settings might be available to overclock the performance (not recommended for servers).

Upgrading Memory

Memory module upgrades will usually be DIMMs. (SIMMs are rarely found on current servers, though older workstations might still use them.) Fortunately, you cannot accidentally install the incorrect memory technology because the memory units have different installation notches, pin count, or length. For example, you cannot install a DIMM in an RDRAM slot. However, you will still need to verify that the memory you are upgrading is of the proper speed and matches other memory already installed in the system.

Technically, you can mix different speeds of memory, although it is not a good practice. For example, you can install 133 MHz SDRAM modules on a 100 MHz bus with existing 100 MHz modules. However, performance of all modules will be limited to 100 MHz. You cannot mix memory speed within a single SIMM memory bank.

Identifying Memory

Identifying existing memory and the memory the motherboard supports is more time-consuming than actually placing the memory modules. Absent the documentation that came with the server, you can identify memory modules by reading the actual chips on the module, counting the chips, and measuring the length. As discussed in Chapter 3, you can read the numbers on the module's memory chips to determine the speed in nanoseconds (ns), which correlates to a manufactured speed. (For your reference, a portion of Table 3-4 appears here as Table 6-2.) For example, if the number ends in –10, then you have 10 ns speed designed for a 66 MHz bus.

Table 6-2 SDRAM speed

Speed in ns	Rated Speed in MHz
15	PC66
10	PC66
8	PC100
7.5	PC133

Not every manufacturer includes these speed markings on their memory chips.

Determine if an error correcting function such as parity or ECC is present in the module. There are several ways to verify this, but the simplest and most foolproof method is to count the number of chips on the module. If the number is evenly divisible by three, then you have either ECC or parity memory. For example, nine chips on the module is evenly divisible by three (9/3 = 3), identifying the module as ECC or parity. If the part number on each chip is the same, then you have ECC, which includes the error correcting function in each chip. If one of the chips has a different part number, then you have parity memory, because the chip that is different is solely responsible for the parity function on behalf of all the memory chips on the module.

Identify a SIMM as a 72-pin module with a single notch in the bottom center, measuring 4.26 inches, or 108.2 mm (see Figure 6-12), and identify a 168-pin DIMM with two notches at the bottom, measuring 5.26 inches, or 133.8 mm (see Figure 6-13). The notches in the bottom are spaced slightly differently. The left notch spacing defines the module as registered, buffered, or unbuffered, and the right notch spacing defines the module

voltage at 5.0V or 3.3V. Again, the notching makes it impossible to make a mistake with buffering or voltages when installing the modules.

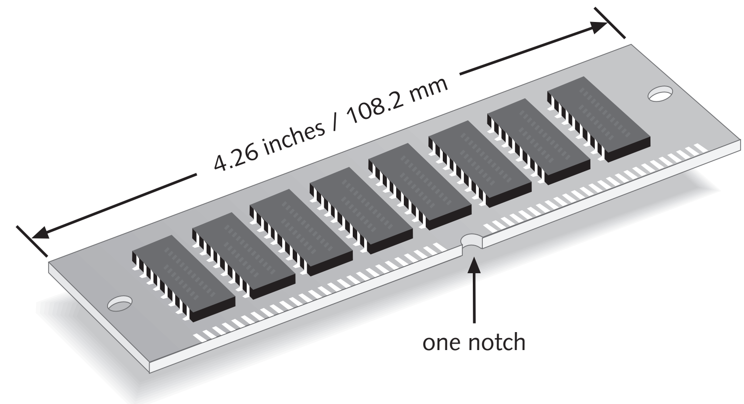

Figure 6-12 A 72-pin SIMM

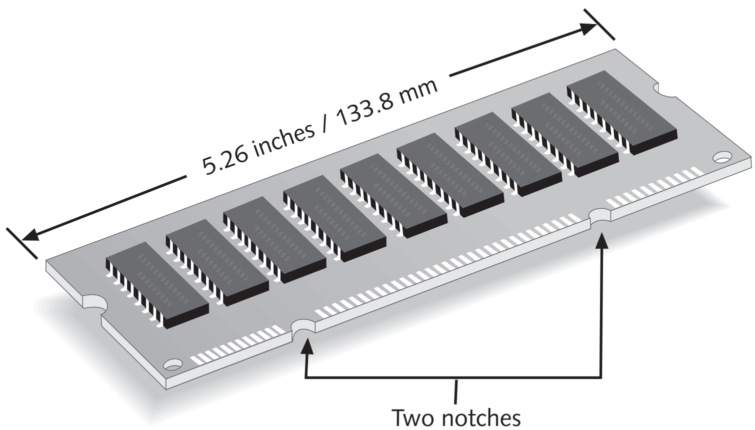

Figure 6-13 A 168-pin DIMM

The next generation of SDRAM, the 184-pin DDR DIMM, is an important memory module with which you should be familiar. High-end workstations and servers are increasingly requiring DDR SDRAM because of its extremely high throughput (up to 2656 MBps). You physically insert this module into the slot in the same way as the 168-pin DIMM; however, you can identify it with a single key notch at the bottom indicating its voltage (2.5V) and with two notches on either end. The module is the same length as a standard SDRAM DIMM, measuring 5.256 inches, or 133.5 mm (see Figure 6-14).

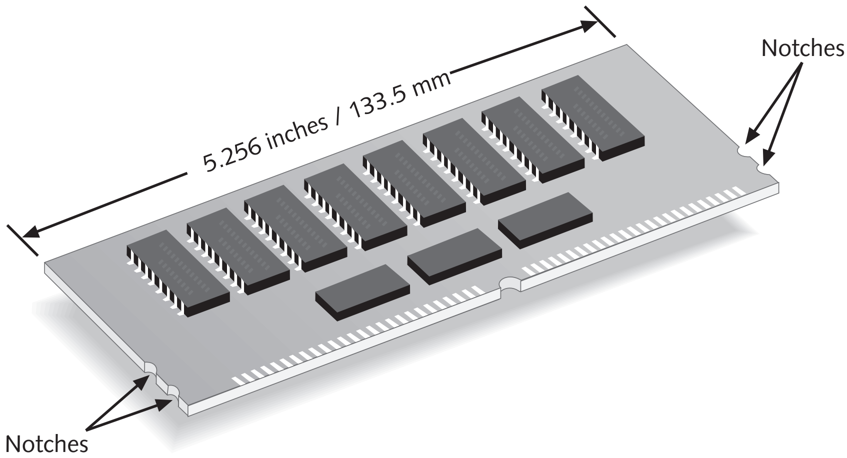

Figure 6-14 A 184-pin DDR SDRAM DIMM

Also look at the physical characteristics of the module. Recall that RDRAM RIMMs have immediately identifiable metal heat spreaders covering the memory chips (see Figure 6-15). Another verification of RDRAM is if empty memory slots on the motherboard have a C-RIMM that completes the continuity, permitting memory data to pass through each slot.

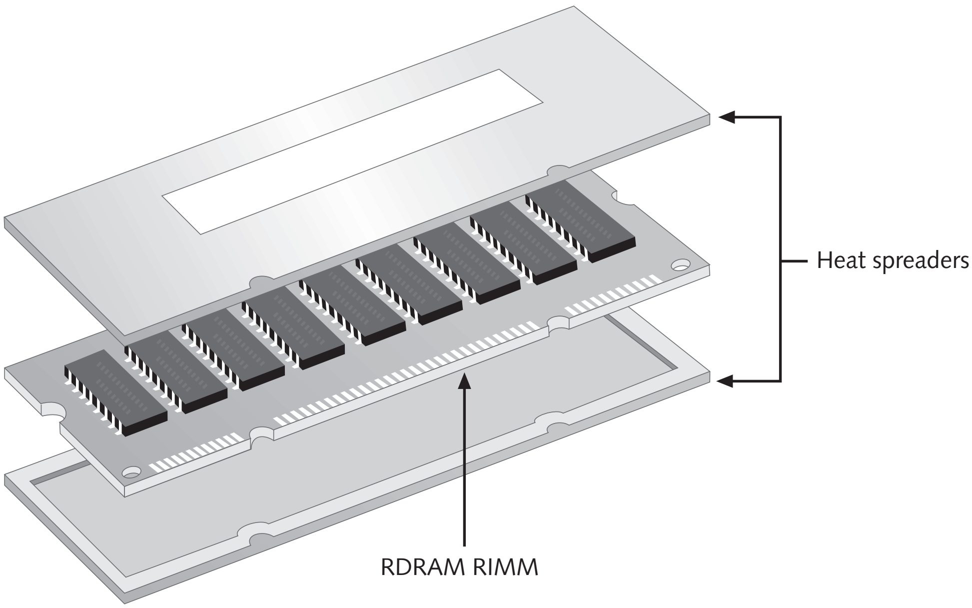

Figure 6-15 An RDRAM RIMM with heat spreaders pulled away

Installing Memory

Installing memory is a straightforward matter. Both SIMMs and DIMMs have notching that prevents you from installing them in the wrong slots. The trick is to make sure that you have fully seated the memory into the slot.

Caution

Memory modules are particularly sensitive to ESD, so be sure to exercise appropriate precautions.

For a SIMM, insert the module at a 45-degree angle to the SIMM slot by pressing down firmly. A notch on one end of the module matches a protrusion on the slot to prevent backward insertion. Tilt the module toward the locking clips until they snap into place. Tabs on the locking clips should fit precisely into holes on the SIMM (see Figure 6-16). If the holes do not appear to match up, you probably have not pressed the module down far enough into the slot. To remove the SIMM, use your fingers to separate the locking clips from the holes in the module while simultaneously tilting the module up and out.

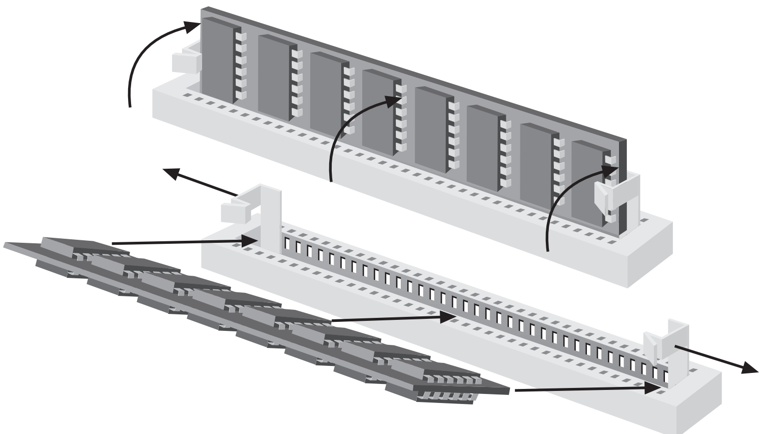

Figure 6-16 Inserting a SIMM into the slot

You are more likely to be installing a DIMM, which requires 90-degree, straight downward insertion into the DIMM slot. However, you will not tilt the module forward like the older SIMM. Instead, locking ejector tabs automatically clamp onto the module when the DIMM is fully seated, though you might also press them into place to verify proper locking (see Figure 6-17). To remove the DIMM, press down on top of both ejector tabs simultaneously, and the module should come out.

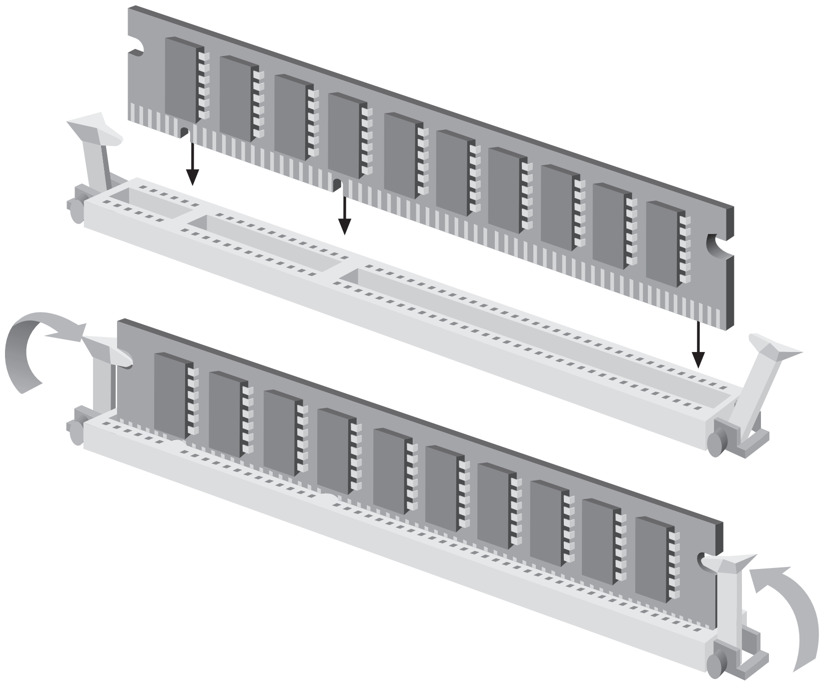

Figure 6-17 The DIMM is inserted with locking ejector tabs

> I've installed DIMMs into extremely tight slots that required considerable pressure to seat them in the slot. When inserting DIMMs, you should feel a confirming "sink" into the slot. Sometimes the ejector tab will flip into place and give the illusion that the module is fully seated when it is not. If you turn on the computer and see memory errors or the system does not execute its POST, check the seating of the modules—they are probably not making contact with the slot.

Remember the following guidelines when installing memory:

- Match the manufacturer of existing memory when possible. Production variances are probably slight, but might produce an unacceptable risk in mission-critical servers. On desktop PCs, mixing manufacturers is a more acceptable risk.
- Try to plan server purchases with as few modules as possible to maximize future upgrade possibilities. For example, instead of two 64 MB DIMMs, install a single 128 MB DIMM.
- If the memory modules vary in size, install the largest modules (in megabytes) in the lowest numbered slot for best performance.
- Verify that the operating system is capable of recognizing the memory you install. For example, Windows NT 4.0 can utilize only 4 GB of memory regardless of motherboard capacity.

- Installing large quantities of memory (2–4 GB or more) might require special configuration of the NOS. Check with the vendor to verify. For example, Windows 2000 Advanced Server allows you to access large amounts of memory by adding the "/PAE" switch (enabling Physical Address Extension) to the Boot.ini bootup file.
- Large quantities of memory might also require specific hardware compatibility. For example, Windows 2000 Advanced Server can utilize PAE only with a Pentium Pro processor or later, 4 GB or more of RAM, and an Intel 450 NX or compatible chipset or later.
- Although the current motherboard and BIOS will automatically detect new memory, verify this in the BIOS to confirm proper seating of the memory.
- Most NOSs include a virtual memory feature that corresponds to the amount of physical memory installed in the system. When you increase the amount of physical memory, also increase the amount of virtual memory, which might be referred to as a swap file. However, in large memory implementations, it might be impractical to create a 4 GB swap file that matches 4 GB of physical memory. In this case, about 2 GB is usually acceptable.

Note

Memory modules and sockets designed for low-cost appeal might use tin contacts. Many sources tell you to avoid mixing gold modules with tin sockets (and vice versa) because contact between the two creates an oxidization known as "fretting corrosion," affecting good electrical contact and possibly creating all kinds of instability and memory error problems. However, this issue is less of a problem than it used to be because tin contacts are rarely produced, especially for server platforms. All industry-standard DIMMs use gold contacts.

Updating the BIOS

In years past, it was impossible to update the BIOS without first removing the BIOS, which might involve soldering tools. Current PCs and servers offer **flash BIOS**, which means you can download the most recent update from the vendor's web site and apply it to the server without replacing the BIOS.

TIP

To keep on top of the most recent updates, consider checking with the server vendor to see if they offer a notification service that sends you an email when a BIOS update (or other updates such as a driver) is available. Dell, for example, offers such a service.

Before updating the BIOS, read all available documentation about the update (a Readme.txt file usually accompanies the BIOS update). This is important to determine the purpose of the update and if it solves any problems you might be experiencing.

Update the BIOS as follows:

1. Download the BIOS update from the system vendor. Although major BIOS manufacturers such as Phoenix Software, Award Software, and American Megatrends Inc. (AMI) make most BIOS found in servers, you should not seek or use updates from the BIOS manufacturer. Server vendors work extensively to tailor a specific BIOS exactly for the server vendor's motherboard.
2. Execute the downloaded file. Typically, this will copy all necessary flash BIOS files to a blank floppy disk.

Caution

The boot disk should be as clean as possible. It should be a system disk containing only basic boot files (IOS.SYS, MSDOS.SYS, COMMAND.COM) but no memory management drivers such as HIMEM.SYS.

3. Record current CMOS settings. Most BIOS updates either automatically reset the CMOS for you or recommend that you manually reset the CMOS to default settings. You can use the recorded CMOS settings to reconfigure the CMOS to your preferences after the update is complete.

TIP

Instead of writing down each setting, consider using the Shift+PrtScn (Print Screen) keys to send the CMOS configuration screen to a locally attached printer. You will have to perform a manual form feed on laser printers to print the page.

4. Boot from the flash BIOS disk. If the extracted files do not create a bootable system disk, use any DOS or Windows 9x system to first format the disk as a system disk.
5. The flash BIOS usually presents a list of options for what you want to do. Select the option to update the BIOS.
6. After the BIOS update is finished, manually reboot the computer if the updated BIOS does not do so automatically.
7. Upon reboot, access the CMOS settings and reset the values back to default. Otherwise, the system might not function correctly.
8. Reboot the server and again access the CMOS settings, entering in your preferred settings recorded in Step 3.

Recovering the BIOS

If the BIOS is corrupt, you cannot boot the system at all—not even with a valid flash BIOS disk. In fact, the display adapter is probably unavailable, further complicating matters. A corrupt BIOS can be caused by a number of things, including ESD to the BIOS EPROM, an interrupted flash BIOS update, or a virus. Regardless of what causes the

corruption, the system cannot function until you repair the BIOS. While some systems might vary (specific instructions are probably in the documentation), the following are the usual steps to take in recovering the BIOS:

1. Turn off the system power.
2. Having removed the server cover, look for a motherboard jumper that allows you to enter recovery mode. Check motherboard documentation to find the jumper if it is not printed obviously on the motherboard.
3. Insert the latest flash BIOS update into the floppy drive.
4. Turn on the system power. A corrupt BIOS usually means that even the video display is not functioning, so you will have to listen to beep(s) to track what is taking place in the BIOS recovery:
   - The BIOS sounds a single beep when it passes control to DOS on the bootable flash BIOS floppy. DOS executes the Autoexec.bat file on the floppy, which in turn runs the flash BIOS update executable.
   - A single beep indicates commencement of the flash operation.
   - Two beeps indicate a recovered system BIOS.
   - Two more beeps indicate successful completion of the recovery.
   - A constant series of beeps indicates a failed recovery attempt.

   The recovery process usually takes between three to five minutes.
5. Remove the recovery diskette, and turn off server power.
6. Restore the motherboard jumper to its original position.

Note

The server might have a read-only jumper position that prevents viruses from writing to the BIOS.

7. Turn on the system power, and access the BIOS settings to enter your preferences.

Note

After rebooting, a CMOS checksum error or some other problem might appear. Try to reboot again (by powering off and on) to see if that resolves the problem. If not, enter the CMOS setup utility to check and save settings. You can often resolve CMOS checksum errors by accessing CMOS settings, saving the settings, and then rebooting.

Upgrading a Power Supply

If you suspect a power supply is not performing reliably or to specifications, you can first verify this using a **digital multimeter (DMM)**. The DMM measures AC voltage, DC

voltage, continuity, or electrical resistance (see Figure 6-18). In this section, we are mostly concerned with the DC voltage readings.

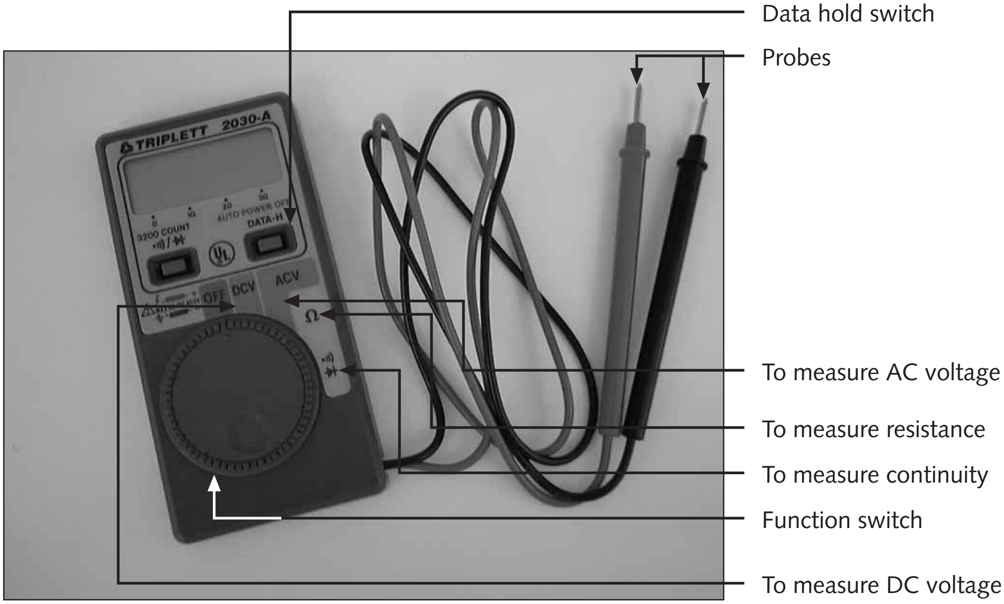

Figure 6-18 A digital multimeter

> **TIP** Memory parity check error messages are often an indicator of a problematic power supply. Recall that memory must be continually charged to retain its data. If the power to memory fluctuates, you are likely to lose data. If parity check messages consistently identify the same location in memory, then the problem is probably with a bad memory module. If the parity check messages are in different locations, then the problem is probably power-related.

Using a Multimeter

A multimeter should help you clarify whether a power supply is operating within specifications. Although a multimeter can be analog (using a needle to show measurements), you should use a digital multimeter for best accuracy.

To check the operating voltage of a power supply using a DMM, use the following procedure:

1. Take the server cover off and power up the system.

2. The DMM has two probes: red and black. Find the power supply connector that connects to the motherboard, and locate the Power_Good pin (pin 8; third pin from the left on the unnotched side). Insert the red probe into the connector at pin 8 (see Figure 6-19). Inserting the probe alongside a live connection like this is known as **backprobing**.

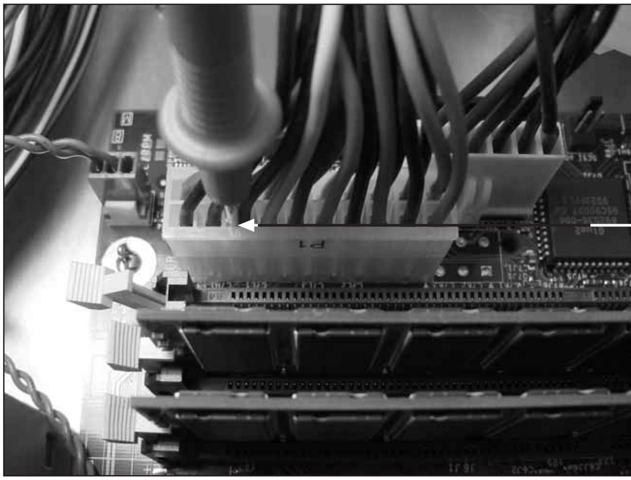

Figure 6-19 Backprobing the ATX power supply connector

3. Touch the black probe to a ground, such as the chassis.
4. The DMM should read between +3 V and +6 V. If not, then the system cannot see the Power_Good signal and does not start or run correctly.
5. Repeat this process for other connectors (use Tables 6-3 and 6-4 as a reference). Connectors are in the +/-3.3 V, +/-5.0 V, and +/-12.0 V range. You should not see more than 10% variance from this range, and only 5% variance is acceptable for high-quality power supplies.

If the voltage readings are outside an acceptable range, replace the power supply.

Table 6-3 Pin Assignments for the ATX Power Connection

Pin Number*	Color	Voltage
1	Orange	+3.3 V
2	Orange	+3.3 V
3	Black	Ground
4	Red	+5 V
5	Black	Ground
6	Red	+5 V
7	Black	Ground
8	Gray	Power_Good
9	Purple	+5 VSB (Standby)
10	Yellow	+12 V
11	Orange (or Brown)	+3.3 V
12	Blue	–12 V
13	Black	Ground
14	Green	PS_On (power supply on)
15	Black	Ground
16	Black	Ground
17	Black	Ground
18	White	–5 V
19	Red	+5 V
20	Red	+5 V

*Pins 1–10 appear on the non-keyed side, and 11–20 on the keyed side.

Table 6-4 Drive Connections

Pin Number	Color	Voltage
1	Yellow	+12 V
2	Black	Ground
3	Black	Ground
4	Red	+5 V

If the DMM requires you to specify a maximum voltage range before testing equipment, set it at 20 V, because servers use +5 V or +12 V. Setting it too low might “peg the meter,” overloading the DMM and possibly damaging it. Many higher-quality DMMs have autoranging capability to automatically determine the best setting.

Replacing the Power Supply

Wise server choices include redundant power supplies, so hopefully replacing a power supply, while urgent, does not involve downing the server. Many redundant power supplies include lights, beeps, or both to alert you of an impending failure as well as total failure. Naturally, you want to replace the power supply when warning of impending failure occurs, instead of waiting for total failure. Check with the vendor documentation for specific instructions on replacing the power supply (PSU). Generally, replacing a hot-swappable redundant power supply is quick and easy:

1. Verify that the replacement PSU is compatible. For hot-swappable server PSUs, I recommend using the server vendor PSU instead of a third-party PSU for best reliability.
2. If necessary, unscrew the power supply from the chassis. Normally, redundant power supplies use a lever or handle instead of screws to attach to the chassis. Flip the lever/handle or unscrew, and then pull the PSU out of the server (see Figure 6-20). You will probably perform this action from the back of the server in most cases, so make sure you have enough space to either slide the server out of the rack far enough to reach behind the server from the front of the rack or that sufficient space exists behind the server to both open a door (if present) and remove the power supply from the back of the rack. Be sure to use both hands when removing the PSU; it might surprise you how heavy it can be.

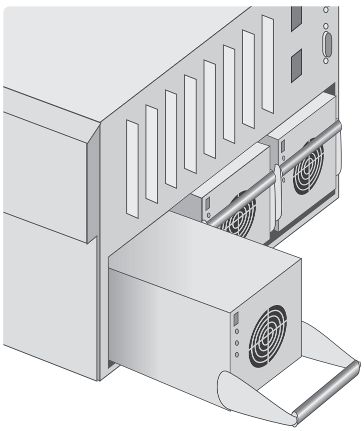

Figure 6-20 Pull down the handle and pull out the hot-swappable PSU

Caution: Verify that the remaining power supply or power supplies are capable of powering the server and all components prior to removing a PSU from the server. Otherwise, the server may go down hard and possibly damage the operating system, applications, or data.

3. Replace the power supply by reversing the steps above. Hot-swappable PSUs plug directly into connections on the server without requiring you to remove the cover and attach power cables individually.

If the replacement power supply is an upgrade, not a matter of replacing a failed PSU, then you can perform the above steps one PSU at a time until all are up to the new level. Be sure to verify vendor documentation to see if the server supports the new power level.

For power supplies that are not hot-swappable, continued operation of the server is not possible when the power supply fails, and replacing the PSU is an immediate concern. Do *not* attempt to repair the PSU; even unplugged it retains a high level of dangerous electricity. It is more prudent to spend the money on a new PSU. Also, if it suits you, there is more flexibility in choosing a different vendor with non-hot-swappable power supplies. Make sure that the new power supply will fit the chassis. Some server vendors use specially designed PSU form factors that prevent you from choosing a generic replacement.

Replace the non-hot-swappable PSU as follows:

1. Power down the system if it is running.
2. Unplug the power cable.
3. Remove the server cover.

Caution

At your discretion, consider *not* grounding yourself in this rare case. If there is a chance that the PSU might ground to you, extremely high voltage could be discharged, resulting in serious injury. Just keep touching the chassis and be extra careful about not touching other ESD-sensitive components.

4. Unplug all PSU cable connections to the motherboard, fans, drives, and so on.
5. Unscrew the PSU from the chassis. Non-hot-swappable PSUs do not have handles or levers to release them from the chassis. Some servers might require that you first remove other components such as adapters or hard disks in order to reach the PSU.
6. Remove the PSU. (If you have time, this is a good time to blow out or vacuum out dust.)
7. Replace with a new PSU by reversing the steps above.

TIP

I don't think you'll see this configuration, but older AT power supply connections use two connectors to connect to the power supply. If you see this, make note of the connector orientation, because mixing them up can burn up the motherboard or components. Also, it's easy to accidentally make the "off by one" connection where you skew the connector to the board by one pin. Remember that the two black ground wires in each connector should be adjacent to avoid a mix-up (see Figure 6-21). ATX-style motherboards, on the other hand, use a single connector that is keyed to prevent backward installation.

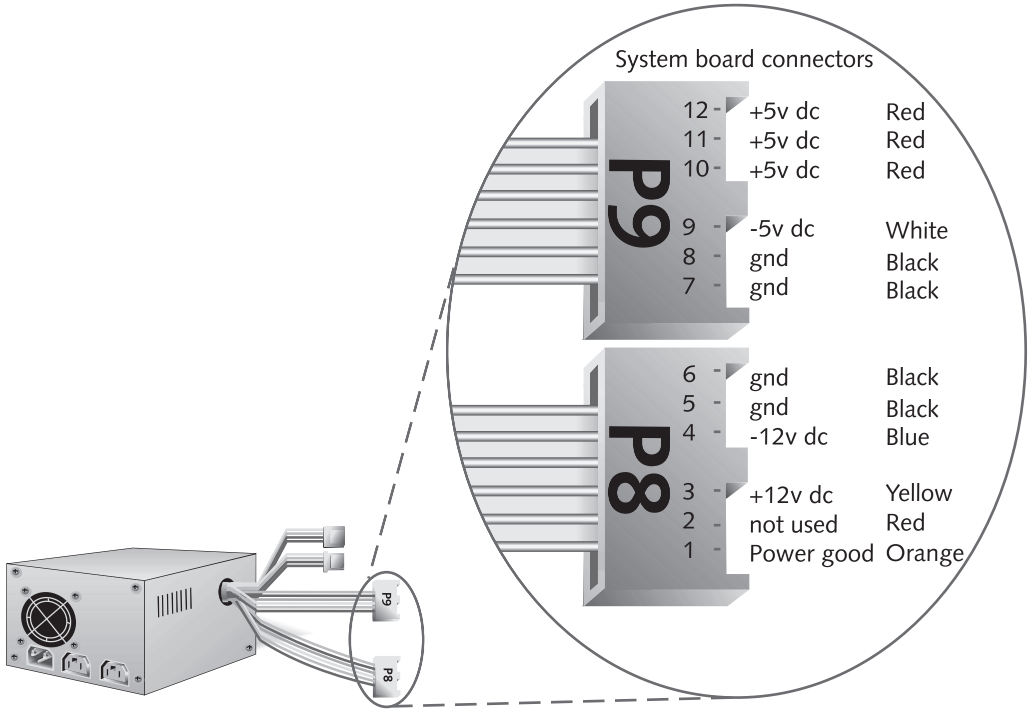

Figure 6-21 Older AT power supply connections

Upgrading Adapters

Upgrading adapters is a fairly straightforward process. Actually, you cannot normally upgrade adapters—any features or characteristics of the adapter are static except for driver updates, BIOS updates (as is the case with SCSI adapters), video adapters with upgradeable memory, or LAN adapters to which you can add a "wake on" LAN chip for network booting. To install new features, simply replace the adapter. When installing the adapter into a slot, the following guidelines define what you can do:

- An 8-bit ISA card can fit into a 16-bit ISA or EISA slot.
- A 16-bit ISA card can fit into a 16-bit ISA or EISA slot.
- A 32-bit PCI card can fit into a 32-bit or 64-bit PCI slot.
- A 64-bit PCI card only fits into a 64-bit PCI slot.

When installing the adapter, apply firm, even pressure when guiding the adapter into the slot. Be sure to orient the server so that you apply downward pressure (toward the tabletop) instead of sideways to avoid tipping the server over. Try to visually line up the card with the slot to ensure success. Often, you will have to remove adjacent cards first for better visibility and more working room. As you press down on the adapter, you should

feel the card "sink" into the slot (see Figure 6-22). If the retaining bracket is not flush against the chassis, then the card is probably not fully seated. Once seated, screw in the retaining screw. Be sure to cover empty slot openings on the back of the chassis to help provide good airflow.

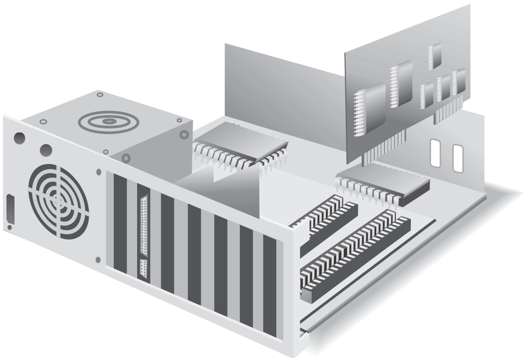

Figure 6-22 Inserting an adapter into its slot

Note

Though unusual with quality components, I have sometimes found cards that could not be fully seated because the retaining bracket was flush against the chassis, preventing further downward motion. If this is the case, you should choose a different adapter. I have also seen cards that, when fully seated, still had clearance between the retaining bracket and chassis. Again, choose a different adapter if possible. Otherwise, you might have to use a pair of needle-nose pliers and carefully bend the adapter to make it fit.

TIP

The AGP graphics slot is further away from the fastening point on the chassis than other bus interfaces, such as PCI. As a result, AGP cards are a little more susceptible to "walking" out of the slot during shipment. When you receive a new server, I recommend pressing down on all removable components, including the AGP adapter, to ensure proper seating of cards. New motherboards often include an AGP card retention mechanism that snaps over the slot and locks the card into place via a retention notch.

Upgrading adapter drivers is normally a matter of connecting to the vendor web site, downloading the drivers, extracting them to a temporary location on the hard drive, and installing them. Hands-on Project 6-7 guides you through this process step-by-step.

Upgrading the UPS

When upgrading the UPS, realize that the upgrade requires you to power down all systems connected to the UPS, because an "upgrade" is really a replacement. For example, if a rack has three 8-way servers in it and a large UPS for redundant power, and the

administrator knows that the UPS will not supply adequate power after adding another 8-way server, he or she might proceed as follows:

1. Power off all load equipment attached to the UPS.
2. Replace the UPS with a model having sufficient power capacity and runtime.
3. Plug in all equipment and again power up the systems on the new UPS.

Of course, because the servers must be offline while the UPS is being replaced, redundant servers must perform the same services. If this redundancy is not available, notify the users of the planned downtime and perform the upgrade during low usage periods. Also, perform a backup prior to the upgrade. The batteries on a new UPS probably do not hold enough charge to provide adequate runtime if utility power were to fail soon after the upgrade, and a recent backup will assist the recovery.

Alternatively, if the rack has N+1 UPS redundancy (as discussed in Chapter 4), then the administrator can upgrade one UPS while another continues to provide redundancy. If utility power is interrupted during the upgrade, the redundant UPS will continue to supply power.

SERVER UPGRADE CHECKLIST

The following checklist should help you to perform carefully planned upgrades:

- ❒ Set a baseline of acceptable performance.
- ❒ Confirm NOS compatibility.
- ❒ Notify users if the upgrade affects them.
- ❒ Record settings if applicable (e.g., CMOS settings prior to BIOS upgrade).
- ❒ Verify available resources for the device (IRQ, DMA, I/O).
- ❒ Download the most recent driver(s) and BIOS upgrade.
- ❒ Read instructions, FAQs, and newsgroups.
- ❒ Inventory delivered parts.
- ❒ Perform a tape backup.
- ❒ Upgrade to the latest BIOS version.
- ❒ Perform the physical installation using sound ESD practices.
- ❒ Test and pilot the implementation.
- ❒ Reset the baseline.

Chapter Summary

- Upgrading a server can require a significant investment in time, planning, and, of course, money—especially for larger upgrades that have a significant impact on the operations of the IT department, users, or your business audience.
- Many times, upgrading the server is a necessary step in response to poor server performance. To fully justify expenditures in time and money, you should create a performance baseline so that you can define an acceptable level of performance.
- You might take that same database server and consider not only upgrading it to an acceptable level of performance, but also upgrading it beyond current needs in a proactive approach to extend the server investment.
- Upgrade the server only during periods of lowest utilization to minimize the impact on users.
- Notifying users helps to reduce the administrator's visibility and avoid unnecessary calls to the IT department. Public advance notice such as an email broadcast documents that you made a reasonable effort to notify users, and tells them how the upgrade will benefit them.
- Before performing the upgrade, apply the latest flash BIOS upgrade and the most recent available drivers.
- Read all documentation and available FAQs or newsgroups relevant to the upgrade and NOS installation.
- The potential impact of some upgrades (both if they succeed and if they fail) might require an isolated pilot program where you can thoroughly test the upgrade for reliability and performance prior to full deployment throughout the organization.
- Perform baseline performance tests before and after the upgrade to measure any changes in performance.
- Log upgrades to an easily accessible location to assist in troubleshooting and asset tracking.
- Always back up the server before performing any hardware or software upgrades.
- You can use clustering or a hot spare to continue providing service during the upgrade process.
- Before upgrading, verify the availability of slots and resources.
- Verify inventory of parts on hand and delivered parts for accounting purposes and to ensure that the upgrade proceeds smoothly.
- ESD poses a threat to electronic devices, and administrators should implement antistatic measures such as an antistatic wrist bracelet and mat.

- When upgrading a processor, verify that the motherboard and BIOS can accept the new processor. If adding a processor for SMP, match the new processor's stepping, cache, form factor, and speed exactly to the existing processor.
- Processors fit in either a socket or slot. Sockets have a ZIF lever to lock the processor into place, and slots have a retention mechanism on the side. Both formats are keyed so that the processor can only be inserted in the correct orientation.
- For Windows NT or Windows 2000, you must take measures to notify the operating system of an additional processor.
- When adding SIMM or DIMM modules, match the speed of existing modules.
- If the number of SIMM/DIMM modules is evenly divisible by three, then you have either ECC or parity memory. If the part number on each chip is the same, then you have ECC, which includes the error correcting function in each chip. If one of the chips has a different part number, then you have parity memory.
- RDRAM RIMMs have metal heat spreaders covering the memory chips.
- For a SIMM, insert the module at a 45-degree angle to the SIMM slot by pressing down firmly. Tilt the module toward the locking clips until they snap into place.
- Installing a DIMM requires 90-degree, straight downward insertion into the DIMM slot. Locking ejector tabs automatically clamp onto the module when the DIMM is fully seated, though you might also press them into place to verify proper locking.
- Match the memory module manufacturer when possible. Try to plan server purchases with as few modules as possible to maximize future upgrade possibilities. If the memory modules vary in size, install the largest modules (in megabytes) in the lowest numbered slot for best performance.
- Verify that the operating system is capable of recognizing the memory you install. For example, Windows NT 4.0 can utilize only about 4 GB regardless of motherboard capacity. Installing large quantities of memory (2–4 GB or more) might require special configuration of the NOS. Large quantities of memory might also require specific hardware compatibility.
- Although current motherboards and BIOS will automatically detect new memory, verify this in the BIOS to confirm proper seating of the memory.
- When you increase the amount of physical memory, also increase the amount of virtual memory and adjust the size of the swap file.
- Current PCs and servers offer flash BIOS, which means that you can download the most recent update from the vendor's web site and apply it to the server without replacing the BIOS. Before you perform a flash BIOS upgrade, be sure to record existing CMOS settings because you must reenter them after the upgrade.
- You can recover a corrupt BIOS using a flash BIOS floppy, but only by using beeping sounds to track the progress because the display is not available.

- If you suspect a power supply is not performing reliably or to specifications, verify PSU performance by using a digital multimeter (DMM).
- Do *not* attempt to repair a PSU; even unplugged it retains a high level of dangerous electricity.
- An 8-bit ISA card can fit into a 16-bit ISA or EISA slot, a 16-bit ISA card can fit into a 16-bit ISA or EISA slot, a 32-bit PCI card can fit into a 32-bit or 64-bit PCI slot, and a 64-bit PCI card only fits into a 64-bit PCI slot.
- Upgrading the UPS is really a replacement, and unless you have redundant UPS systems powering the rack, you must plan for downtime.

Key Terms

backprobing — Inserting the probe alongside the live connection.

baseline — Performance data that reflects an acceptable level of system performance.

digital multimeter (DMM) — A device that measures AC voltage, DC voltage, continuity, or electrical resistance.

Direct Memory Address (DMA) — A resource that ISA devices use to directly access memory without first having to access the processor, both increasing device performance and reducing processor load. There are eight DMA channels, numbered 0–7.

driver — A software interface that allows the hardware to function with the operating system.

electrostatic discharge (ESD) — A discharge of electrical energy that occurs when two objects with differing electrical potential come into contact with one another because the electrical charges seek to equalize.

flash BIOS — BIOS that can be updated via software instead of physical removal of the EPROM.

heat sink — An attachment to the processor that either dissipates heat through cooling fins or a small cooling fan in addition to cooling fins.

hot spare — A specific component (usually a hard drive) or a complete server that can immediately perform on the network and transparently perform the exact same function as the original.

interrupt request (IRQ) — Request that a device uses to "interrupt" the processor to ask for processor resources. There are 16 IRQs, numbered 0–15.

I/O port — A location in memory that the processor uses to communicate with the device.

memory address — Some devices reserve a dedicated region in system memory that is unavailable for use by any other device, application, or the operating system. This can help device stability by ensuring that nothing else trespasses the memory, which causes system errors.

pilot program — Isolating an upgraded server in a portion of the network that makes performance determination easier to determine and lessens negative impact should some part of a major upgrade fail.

Pin Grid Array (PGA) — An arrangement of pins on the underside of a processor. The pins fit inside a corresponding PGA socket.

S-spec — An alphanumeric code printed on the processor that uniquely identifies the processor version and is more specific than processor stepping.

Single Edge Contact Cartridge (SECC) — A slot format processor that stands upright inside a motherboard slot, similar to adapter or memory slots.

Single Edge Contact Cartridge2 (SECC2) — A longer form of the SECC slot that accommodates Pentium Xeon processors.

Staggered Pin Grid Array (SPGA) — Same as a PGA processor or socket format, except in a staggered arrangement to squeeze more pins in the same space (as opposed to straight rows).

stepping — The version of a processor.

zero insertion force (ZIF) — A socket format that allows gravity alone to seat the processor. The processor is then locked into place with a locking lever.

Review Questions

1. Which of the following might be cause to upgrade the server? (Choose all that apply.)
   a. You want the latest technology.
   b. Server performance is insufficient.
   c. timing in relation to the budget
   d. planning for future events
2. When should you upgrade the server?
   a. during peak utilization
   b. during lowest utilization
   c. between 12:00 AM and 3:00 AM
   d. during the weekend
3. Why should you notify users of a planned upgrade? (Choose two.)
   a. to inform users as to how the upgrade benefits them
   b. to increase visibility of the administrator so that people know you're actually working
   c. to reduce support calls during server downtime
   d. to let users know who to call if the upgrade seems to be taking too long

4. Which of the following is an important factor in ensuring that the hardware operates correctly with the operating system?
   a. BIOS
   b. parity memory
   c. error correcting code (ECC)
   d. bus speed
5. If the cover is off the server during the upgrade, you might as well:
   a. leave it off for optimum cooling
   b. clean the dust from the inside
   c. test each electrical lead
   d. upgrade the memory
6. Where can you check to see if an upgrade might cause a problem? (Choose all that apply.)
   a. vendor's web site
   b. FAQs
   c. newsgroups
   d. all of the above
7. A pilot program ensures which of the following? (Choose two.)
   a. safety for passengers
   b. that the upgrade is reliable
   c. that an upgrade failure affects a smaller scope
   d. that the upgrade is successfully deployed across the organization simultaneously
8. In a real production enterprise, you should use a pilot program for which of the following? (Choose two.)
   a. changing any hardware, no matter how small
   b. upgrading to a different NOS
   c. changing the processor platform
   d. upgrading the monitor to a larger size
9. The purpose of the log book is to:
   a. blame someone else if a problem occurs
   b. assist in troubleshooting upgrades
   c. record system errors
   d. record when tape backups are performed

10. Before an upgrade, why would you perform a full backup instead of depending upon the normal tape rotation?
    a. The normal tape rotation might not be as current.
    b. The tape in the rotation might not have sufficient remaining space.
    c. The normal tape rotation might not backup all the data.
    d. You want the backup to reflect the most recent changes to the data.
11. Which of the following is a resource you should verify is available prior to a hardware upgrade?
    a. IRQ
    b. serial port
    c. parallel port
    d. electrical load
12. If you touch a server component and discharge static, it is only a problem if you can:
    a. see the spark
    b. feel the electrical discharge
    c. hear the electrical discharge
    d. always damage components with static discharge
13. Absent any grounding equipment, what can you do to protect against ESD?
    a. touch the server power cable
    b. touch the server chassis
    c. discharge ESD against something else first
    d. leave the server plugged in
14. What is the purpose of taking inventory? (Choose two.)
    a. to fairly distribute server equipment among all sites in the enterprise
    b. to assist in budgetary projections
    c. to avoid having to halt the upgrade because of missing parts
    d. none of the above
15. Which of the following is not important in upgrading the processor?
    a. verifying stepping, cache size, and speed
    b. verifying that the new processor is compatible with the motherboard
    c. verifying that the new processor is compatible with the BIOS
    d. verifying that the processor is at least 400 MHz

16. Into which of the following socket or slot can you install a Pentium III Xeon 700 MHz processor?
    a. Socket 370
    b. Socket 490
    c. Slot 2
    d. Socket A
17. What can you identify about a memory module with nine chips, one having a different part number on it than the others?
    a. The memory is a RIMM.
    b. The memory is EDO memory.
    c. The memory is ECC.
    d. There is insufficient data to make a determination.
18. What should you do prior to a flash BIOS upgrade?
    a. record all BIOS settings
    b. reset BIOS to default settings
    c. remove the CMOS from the motherboard
    d. remove the CMOS battery
19. Which of the following does a multimeter not test?
    a. resistance
    b. voltage
    c. continuity
    d. MHz
20. Which of the following applies to a hot-swappable power supply? (Choose all that apply.)
    a. It is removable with a handle or lever.
    b. It screws onto the chassis.
    c. It is not technically hot-swappable; you must first plug/unplug all power cables.
    d. It does not supply redundancy unless you have an N+1 configuration.

HANDS-ON PROJECTS

Project 6-1

Hands-on Project

In this project, you will update the flash BIOS. It does not matter if the system already has the current version of the BIOS; you can still flash it.

1. Turn on the server and check to see what version of the BIOS is on the system. If the screen flashes by too quickly, you can usually freeze it by pressing the **Pause** button on the keyboard, and then resume by pressing the **Spacebar**. Make a note of the BIOS version.
2. Access the BIOS settings, and record whatever settings you see that look like settings you might want to restore later. For purposes of this course, it is probably OK to bypass restoring the settings later because the server is not an in-use production server and it is not important that preferences and configuration be restored.
3. Connect to the server vendor's web site (not the web site of the BIOS manufacturer).
4. Navigate to the support pages, and locate the downloads for the server. Download the most recent BIOS update.
5. Format a floppy and leave it in the drive.
6. Execute the BIOS download by double-clicking it from the vendor's web site. Follow the on-screen instructions. When finished, the flash BIOS should be on the floppy.
7. Reboot the server with the flash BIOS floppy in the floppy drive. If the system boots to the NOS, then properly shut down the system, reboot again, and adjust the BIOS so that it boots from the floppy before the hard disk. If the system cannot boot from the floppy, you must make the floppy bootable. Do this by accessing a Windows 9x computer and from an MS-DOS prompt, type **SYS A**:. The necessary system files then transfer to the floppy.
8. The flash BIOS offers you an option to perform the flash BIOS update; go ahead and do so.
9. If the system does not reboot at the end of the flash BIOS update, reboot it.
10. View the BIOS version upon reboot. Is it a newer version?
11. Access the CMOS settings. Have any customized settings been reset to default settings?
12. Reset the CMOS to defaults, save the settings, and reboot.
13. Confirm a successful boot to the operating system.

Project 6-2

In this project, you will determine available system resources on a Windows 2000 server.

1. Start up and log on to a Windows 2000 server using the Administrator account.
2. Right-click the **My Computer** icon on the desktop and click **Manage**. The Computer Management screen appears.
3. Click **Device Manager** in the left tree pane.
4. Browse various devices in the right pane by expanding the device category and double-clicking specific devices. For example, double-click the **Network adapters** category, and then double-click a specific network adapter.

5. Select the **Resources** tab of a device.
6. What Input/Output Range (that is, the I/O port), Memory Range, DMA, and Interrupt Request does the device use? (Be careful *not* to make changes.) Click cancel to close the network adapter's properties.
7. Simplify the view so that you don't see devices as much as you see the resources in use. Click the **View** menu and then select **Resources by type**. Now, you should see all four categories of resources. Browse through the resources. In the Interrupt Request category, you might see several devices using the same IRQ (probably IRQ 9). Recall from Chapter 3 that this is because of IRQ steering, which allows the sharing of an IRQ among PCI devices.
8. If a printer is available, print out a summary by clicking the **View** menu, and then **Print**.
9. Close the Computer Management window.

Project 6-3

In this project, you will connect an ESD grounding kit to the server chassis in preparation for the other projects to follow.

1. Unplug the power cord from the server.
2. Remove the cover from the server.
3. Attach the wrist strap to your wrist. If necessary, adjust the wrist strap first.
4. Snap the cord to the wrist strap if not already connected.
5. Snap the other end of the cord to the grounding mat.
6. Connect the cord from the grounding mat to the server chassis using the alligator clip.
7. If the mat has another cord that connects to the ground prong of a power cord plugged into an electrical outlet, go ahead and connect it.
8. You are now at equal electrical potential to the server, and should not discharge any static when touching server components.

Project 6-4

In this project, you will use a multimeter to check the power supply power using backprobing.

1. Using a digital multimeter, set the controls to read 20 V DC.
2. Plug in the server and turn on the power. Leave the cover off.
3. It might be best to tip the server on its side if it's a tower, so that the motherboard is parallel with the tabletop.
4. Find the power supply connector that attaches to the motherboard, and locate pin 8 (use Table 6-3 as a guide).

5. Insert the red probe into the connector along pin 8 so that it touches the metal connection inside the connector.
6. Touch the black probe to the chassis for grounding.
7. Read the multimeter. It should read between +3 V and +6 V. What is the voltage reading?
8. Repeat this process for the remainder of the motherboard connector pins, again using Table 6-3 as a guide. Approximately what percentage variance do you typically see?
9. Backprobe the +5 V and +12 V pins on a hard disk connector (use Table 6-4 as a guide).
10. Leave the server cover off.

Project 6-5

In this project, you will add a processor to a motherboard with an existing processor.

1. Shut down and power off the server if it is still on. The server should have Windows 2000 Server or Advanced Server installed using a single processor.
2. Verify that your grounding kit and wrist strap are connected, and unplug the power cord from the server.
3. Identify the type of socket or slot the motherboard accepts, and the kind of processor that is installed in the server. If you cannot easily identify this information, power up the system and access the BIOS, where you should be able to find the information. Write down the information on a separate piece of paper.
4. With the power off, install a matching processor into the other slot or socket.
5. Power up the system again. Although Windows 2000 Server runs fine, it is not using the second processor until you change the Computer driver to ACPI Multiprocessor PC or MPS Multiprocessor PC as follows:
   a. Right-click **My Computer** and choose **Manage**.
   b. Select **Device Manager** in the left tree pane.
   c. Expand the Computer node in the right tree pane. The type of computer appears under Computer, and is probably Standard PC or Advanced Configuration and Power Interface (ACPI) PC.
   d. Whichever type of computer appears, right-click it and choose **Properties**.
   e. Click the **Driver** tab of the Properties sheet.
   f. Click the **Update Driver** button.
   g. Click **Next** to skip the introduction of the Upgrade Device Driver Wizard.
   h. Select **Display a list of the known drivers for this device so that I can choose a specific driver**, and then click **Next**.

i. Select the **Show all hardware of this device class** option. The screen shown earlier in Figure 6-11 appears.

j. Select the multiprocessor option that matches your computer, and proceed to the end of the wizard.

6. Follow the prompt that appears requesting you to reboot the server.
7. After the reboot, click **Start**, point to Programs, point to Administrative Tools, and click **Performance**. On the toolbar, click the **Add** button to display the Add Counters dialog box. By default, it opens to the Processor performance object. In the right side of the interface under *Select instances from the list*, you should see _Total (aggregate of all installed processors), 0 (the first processor), and 1 (the second processor). This verifies that Windows 2000 Server properly detected both processors.
8. Click the **All instances** radio button and close the Add Counters dialog box.
9. Perform a number of activities (open a program, play a sound clip, and so on). In the Performance window, you should see activity from both processors.
10. Close all open windows and shut down the server.

Project 6-6

Hands-on Project

In this project, you will add memory to the motherboard.

1. With the case still open from the previous project, look at the memory module(s) installed on the motherboard. Using the information presented in this chapter, identify what kind of memory it is—SIMM, DIMM, or RIMM; include information on parity, ECC, and speed where relevant.
2. Remove the memory module and reinstall it.
3. Start the server. If the system does not POST, the module is not fully seated in the socket. Turn off the power and try again.

Project 6-7

Hands-on Project

In this project, you will install a PCI network card using available drivers, and then download the latest drivers from the vendor web site and update the drivers.

1. Remain properly grounded and unplug the server. Install a PCI network card in any available PCI slot on the Windows 2000 server.
2. Start the server. Windows 2000 Server includes many drivers for various NICs and might automatically install drivers. If not, supply drivers from a floppy that accompanies the NIC as prompted by the operating system.
3. Access the network card in Device Manager to verify that it is working properly.

a. Right-click **My Computer**.

b. Select **Manage**.

c. Select **Device Manager** in the left pane.

d. In the right pane, expand the Network adapters item by clicking the "+" sign.

e. Right-click the network card and select **Properties**.

4. The vendor might have provided an updated driver for the network card. Download the newest driver from the Internet (if necessary, use a different computer).
5. Read the vendor instructions that might specify the purpose of the driver update and installation information. In most cases, the driver files are in a self-extracting executable file. Extract the files.
6. Open the NIC property sheet in Device Manager (if it is not already open), and select the **Driver** tab.
7. Click the **Update Driver** button, and click **Next** to start the Upgrade Device Driver Wizard.
8. Select the item **Display a list of the known drivers for this device so that I can choose a specific driver** and click **Next**.
9. Click the **Have Disk** button and navigate to the location of the extracted drivers. Then, finish the wizard.
10. Close all open windows and shut down the server.

CASE PROJECTS

Case Project

1. Your organization of 80 users requires a memory upgrade of the only logon server. The organization keeps typical business hours, but some people start work early or stay late. You have a regular backup that takes place every night at 11:00 PM. What strategy should you use to upgrade this server?
2. Utility power to your server rack of three servers fails. You have a single UPS, but as soon as it starts to supply power, it issues a message that the UPS is overloaded and fails to supply power. Utility power comes back on in a few minutes, and the servers start to boot. What should you do?